UDC 621. 376. 56
621. 397.

RESEARCH DEPARTMENT

1969
THE QUEEN'S AWARD
TO INDUSTRY

REPORT

Pulse code modulation of video signals: 8-bit coder and decoder

No 1970/25

Research Department, Engineering Division
THE BRITISH BROADCASTING CORPORATION

RESEARCH DEPARTMENT

PULSE CODE MODULATION OF VIDEO SIGNALS : 8-BIT CODER AND DECODER

Research Department Report No. 1970/25
UDC 621.376.56
621.397

It uses SI units in accordance with B.S. document PD 5686.

Work covered by this report was undertaken by the BBC Research Department for the BBC and the ITA

D. Maurice

V.G. Devereux, M.A.

Head of Research Department

(EL-42)

Research Department Report No. 1970/25

PULSE CODE MODULATION OF VIDEO SIGNALS : 8-BIT CODER AND DECODER

(EL-42)

July 1970

Research Department Report No. 1970/25
UDC 621.376.56
621.397

PULSE CODE MODULATION OF VIDEO SIGNALS : 8-BIT CODER AND DECODER

Summary

Equipment is described which has been constructed for converting analogue television signals into 8-bit p.c.m. signals and vice versa. The conversion rate used for video signals is approximately 13 x 10^6 samples per second.

Analogue-to-digital conversion is accomplished by two consecutive 4-bit parallel converters. The design of the digital-to-analogue converter is more conventional and employs a binary ladder network to weight the digits in the required manner.

The extensive use of readily available integrated circuits has resulted in equipment which is compact, reliable and relatively cheap to manufacture.

The peak-to-peak signal to quantising noise ratio caused by the combined coding and decoding processes was found to be 57 dB as opposed to 59 dB for a perfect 8-bit system.

1. Introduction

Equipment for converting analogue video signals into 6-bit p.c.m. signals and vice versa has been described in a previous report.[1] Subjective tests using this equipment indicated that seven or eight bits would be required to give a picture quality of the standard required for broadcasting. It was therefore decided that an 8-bit a.d.c. (analogue-to-digital converter) and d.a.c. (digital-to-analogue converter) were needed for further investigations into high quality video digital systems. This report describes the design and performance of such equipment.

The basic principles of operation of the new a.d.c. and d.a.c. are similar to those of the previous equipment, but the circuits have been almost completely redesigned. The main difference is that integrated circuits have been used extensively in place of discrete components, and as a result a considerable reduction in size has been achieved.

The system developed so far is intended for work in which the a.d.c. and d.a.c. are both situated in the same laboratory. These conditions allow simplifications to be made compared to a situation in which the a.d.c. is remote from the d.a.c. As an example, the same clock pulse generator is used to control the operation of both the a.d.c. and d.a.c. Also, the digital signals are handled in eight parallel bit lines; over long distances they would normally be transmitted serially.

2. Analogue-to-digital converter

2.1. General

The new 8-bit a.d.c. uses the same type of parallel/series conversion system as the previous 6-bit video a.d.c. One of the advantages of this system compared with most others suitable for converting video signals[2,3,4,5] is that the speed of operation demanded of the circuits is relatively slow. This feature allows a high proportion of the circuitry to be constructed from readily available and inexpensive integrated circuits.

The conversion rate used for 625-line video signals with a bandwidth of 5·5 MHz is about 13 x 10^6 samples per second. If required, the sampling rate can be varied over a wide range; the prototype unit operated satisfactorily from 30 x 10^3 up to 16 x 10^6 samples per second. The lower limit of this range could easily be reduced if necessary by increasing the time constants of a few a.c. coupling networks, but faster-acting integrated circuits would be required to increase the upper limit.

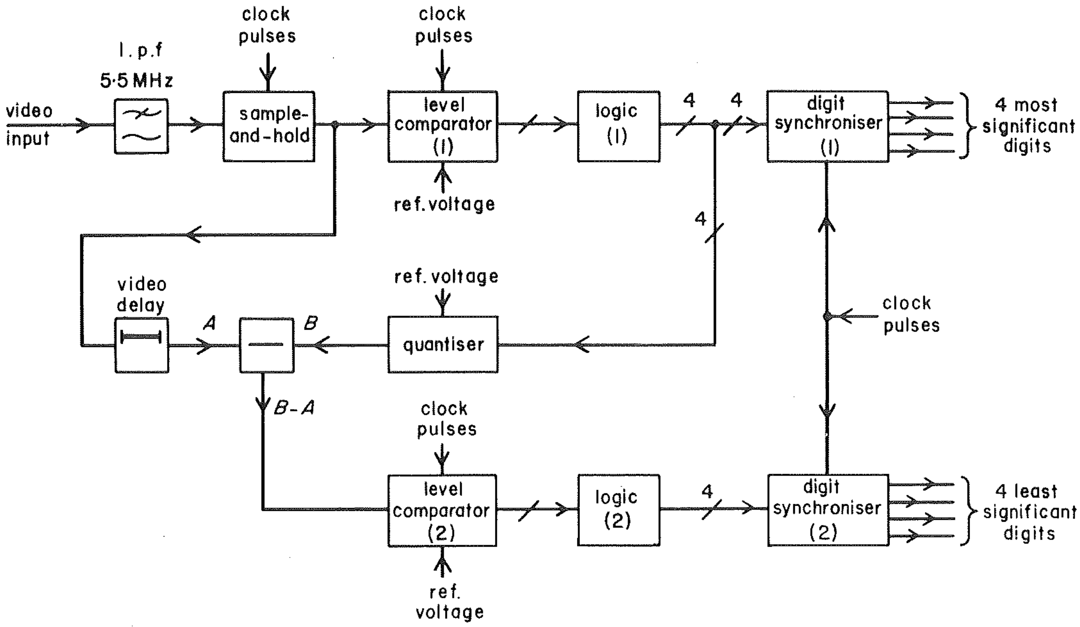

Fig. 1 - Block diagram of video a.d.c.

As shown in Fig. 1, the a.d.c. has two similar stages each of which is a 4-bit parallel converter. The first stage determines the state (1 or 0) of the 4 most significant bits and the second determines the state of the 4 least significant bits.

In the first stage, the sampled and held video signal is fed to 15 level comparators connected in parallel. The outputs of these comparators change state when the video signal crosses voltage levels of (5/16) x 1,2,3 ... 15 volts above earth potential.

A simple logic network converts the signals from the comparators into a 4-digit binary number representing the magnitude of the video signal.

To obtain the lower-order bits, it is first necessary to decode the 4 most significant bits to obtain a 16-level quantised version of the video signal. This quantised signal is then subtracted from a suitably delayed version of the sampled and held video signal. The difference signal thus obtained (see Fig. 2) indicates the amount by which the video signal is greater than the level given by the first 4 bits. This difference signal is divided into a further 16 parts by a second set of level comparators, the outputs of which are decoded to give the state of the 4 least significant bits. The maximum peak-to-peak magnitude of the difference signal is 5/16 volts.

Finally, the digital signals are passed through digit synchroniser units which re-time them so that all 8 digits corresponding to a given sample appear simultaneously at the output of the a.d.c.

This type of converter uses a production-line technique in which the first stage starts operating on a new sample which previous samples are still undergoing conversion in later stages; this is made possible by controlling the timing of successive stages by increasingly delayed clock pulses. It is thus possible for the total conversion time to be greater than the interval between samples; at the sampling rate used for video signals, the conversion time is 240 ns while the sampling interval is only 75 ns. The main advantage of this technique is that each section of the converter is allowed the maximum possible settling time.

Before the relative timing of the various operations is explained, it is convenient to consider some of the circuitry in more detail.

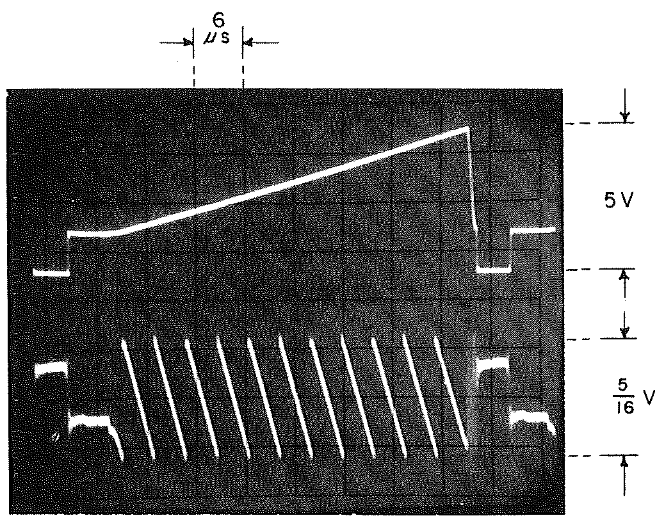

Fig. 2 - Difference signal (lower trace) obtained with line frequency sawtooth applied to input of a.d.c. (upper trace)

2.2. Sample and hold unit

The sample and hold circuit is very similar to that described in Reference 1, but for convenience the important details and modifications are described below. A full analysis of this type of circuit is given in reference 6.

The basic sample-and-hold circuit consists of a four-diode gate and storage capacitor C as shown in Fig. 3. The duration of the sampling pulse is about 25 ns and, for video signals, each sample is held for 50 ns. The diodes used in the sampling gate are of the hot carrier type. Resistance R is adjusted to critically damp the resonant circuit consisting of the storage capacitor and the stray inductances in its charging circuit. The resulting time constant of R and C is about 2·5 ns.

The behaviour of this circuit is illustrated in Fig. 4. During the 25 ns sampling pulse, the voltage on the storage capacitor changes very rapidly until it reaches the level of the video input signal. Thereafter it changes less rapidly as it follows variations in the video signal and is finally held constant at the video level occurring at the instant at which the gate is cut off.

When the black level stabiliser (see Fig. 3) is in use either the bottom of synchronising pulse or black level may be held at the negative end of the conversion range. These alternative clamping conditions are provided so that either the composite video signal or the video signal alone without synchronising pulses may be encoded.

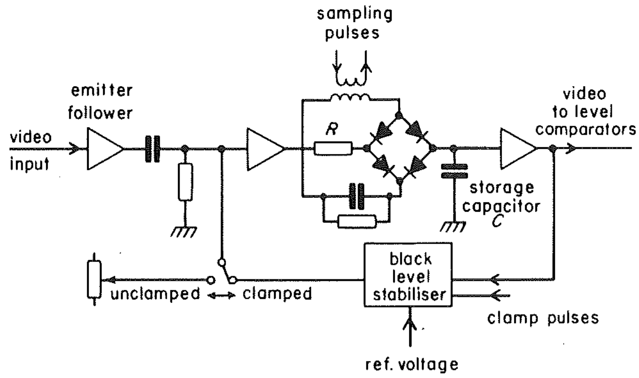

Fig. 3 - Sample and hold unit

Alternatively, the black level stabiliser may be disconnected and the mean d.c. level of the signal is then adjusted by a potentiometer. This facility allows the a.d.c. to be used for signals other than video.

2.3. Level comparators

The output signal from a level comparator has two stable states which indicate whether the applied video signal is above or below a reference potential. Since the peak-to-peak magnitude of the video signal is 5 volts and it is divided into 256 parts, errors in change-over level must be considerably less than 20 mV. A second requirement of the comparators is that the change in their propagation delay with varying amounts of overdrive should be shorter than the time for which the video signal is held constant between samples. Only variations in propagation delay are of importance since a constant delay can be compensated for by delay in the video path. (The effect of delays are considered in more detail in Section 2.7.)

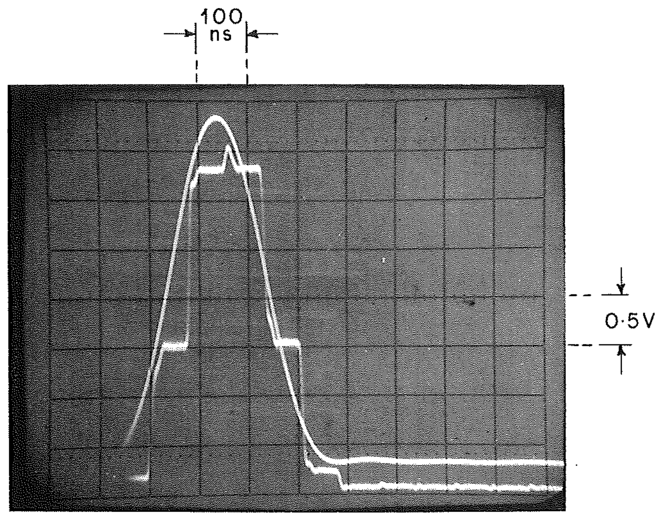

Fig. 4 - Sin² pulse at input and output of sample and hold circuit

A block diagram of a level comparator is shown in Fig. 5. A voltage difference of about 2 mV between the inputs of the differential amplifier causes its output to enter a high or low limiting state. Since it is possible for this amplifier to give an ambiguous output mid-way between its high and low states, it is followed by a D-type flip-flop which samples the output signal and makes a definite decision as to which state it is in. A second function of the flip-flop is to ensure that the output of the comparator cannot change state during the interval between clock pulses.

In practice, it was found to be convenient to use integrated circuits containing pairs of differential amplifiers having their outputs internally connected. The use of such circuits is explained in Appendix 1.

2.4. Logic

The logic system which translates the outputs of the level comparators into a 4-digit binary number is explained in detail in Appendix 1.

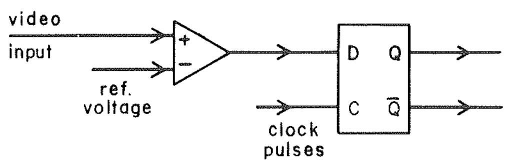

Fig. 5 - Level comparator unit

The only feature which affects further explanation of the operation of the a.d.c. is that all signal paths pass through two NAND gates and therefore the delays thereby introduced into all these paths are equal. TTL integrated circuits are used throughout this section of the a.d.c.

2.5. Subtractor and quantiser

In practice the quantiser and subtractor units are combined in a single circuit, the main details of which are shown in Fig. 6. A 16-level quantised sample is obtained by adding binary weighted currents corresponding to the four most significant digits. These weighted currents are supplied via switches S1, S2, S3 and S4 which are operated by the digital signals in such a manner that the quantised signal obtained is of opposite polarity to the unquantised video signal. The resulting output signal from the amplifier therefore represents the difference between the quantised and unquantised video signals.

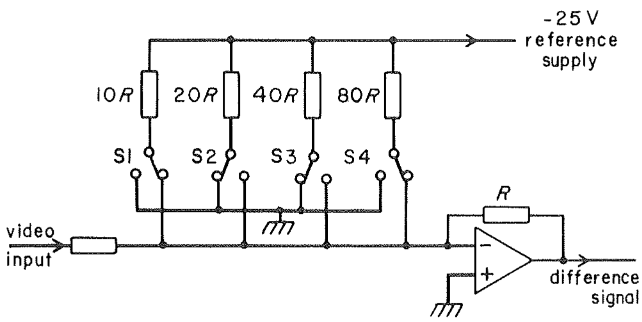

Fig. 6 - Subtractor and quantised signal generator

Special attention was paid to the design of the subtracting amplifier as the following requirements had to be fulfilled.

First, a low input impedance was required to prevent the quantised/unquantised signals interacting with one another.

Secondly, the time taken for the amplifier to settle to within a few millivolts of its final level after a step change in an input signal had to be less than about 30 ns.

Finally, the drift in the output of the amplifier under closed loop conditions should not be more than about 2 mV over the working temperature range.

No suitable integrated circuit amplifier could be found to fulfil all these requirements and it was therefore necessary to use discrete components. A suitable design was achieved using a differential pair gain stage followed by two emitter followers. The open loop gain of this amplifier was 50 dB at low frequencies falling to 0 dB at about 60 MHz.

2.6. Digit synchronisers

These synchronisers (shown in Fig. 1) re-time the two groups of 4 digits from the first and second stages of conversion so that all eight digits from a given sample appear simultaneously at the output.

Since the 4 most significant digits are determined before the 4 least significant digits, digit synchroniser 1 must provide more delay than digit synchroniser 2. This extra delay is obtained by passing each of the 4 most significant digits through a two-stage shift register consisting of two D-type flip-flops in cascade whereas the 4 least significant digits are only passed through a single D-type flip-flop. The relative timing of the clock pulses and digital inputs to these synchronisers is explained in the following section.

2.7. Timing sequence

The relative times at which samples are processed by each stage of the converter are illustrated in Fig. 7. The portions of the waveforms indicated by solid lines correspond to the conversion of two successive samples resulting from the sampling pulses P_1^1 and P_2^1 indicated on waveform (*b*). The timings of these pulses at other points in the a.d.c. are indicated by waveforms (*d*), (*g*) and (*i*), P_1^{11}, P_1^{111} and P_1^{1111} being delayed versions of P_1^1.

It should be noted that clock pulse delays greater than the clock pulse period are used where this would not appear to be necessary. For example, P_1^{11} in waveform (*d*) is derived from P_1^1 rather than from P_2^1 in waveform (*b*). The reason for this is to enable the clock pulse frequency to be varied over a wide range without requiring any alteration in the relative timings of the different clock pulse trains.

The D-type flip-flops in the level comparators and digit synchronisers are triggered by the positive or upward going edges of the clock pulse waveforms.

The triggering edges of the clock pulses to the level comparators are timed to occur at the end of the held portions of the signals applied to them. (See waveforms (*c*) and (*d*).) This gives the differential amplifiers in the comparators the maximum possible time in which to settle before their output states are registered by the following flip-flop. With this timing of clock pulses to the level comparators, a typical output signal from logic (1) appears during the time indicated by waveform (*e*).

In order that the video and quantised signals from a given sample should occur simultaneously at the subtractor, the video signal must be delayed by the interval $(t_1 - t_0)$ between waveforms (*c*) and (*e*). The resulting output signal from the subtractor is of the form indicated by waveform (*f*).

The delays occuring in the second stage level comparators and logic (see waveforms (*f*), (*g*) and (*h*)) are identical to those in the first stage and therefore the digital outputs from logic (2) occur at the time indicated by waveform (*h*).

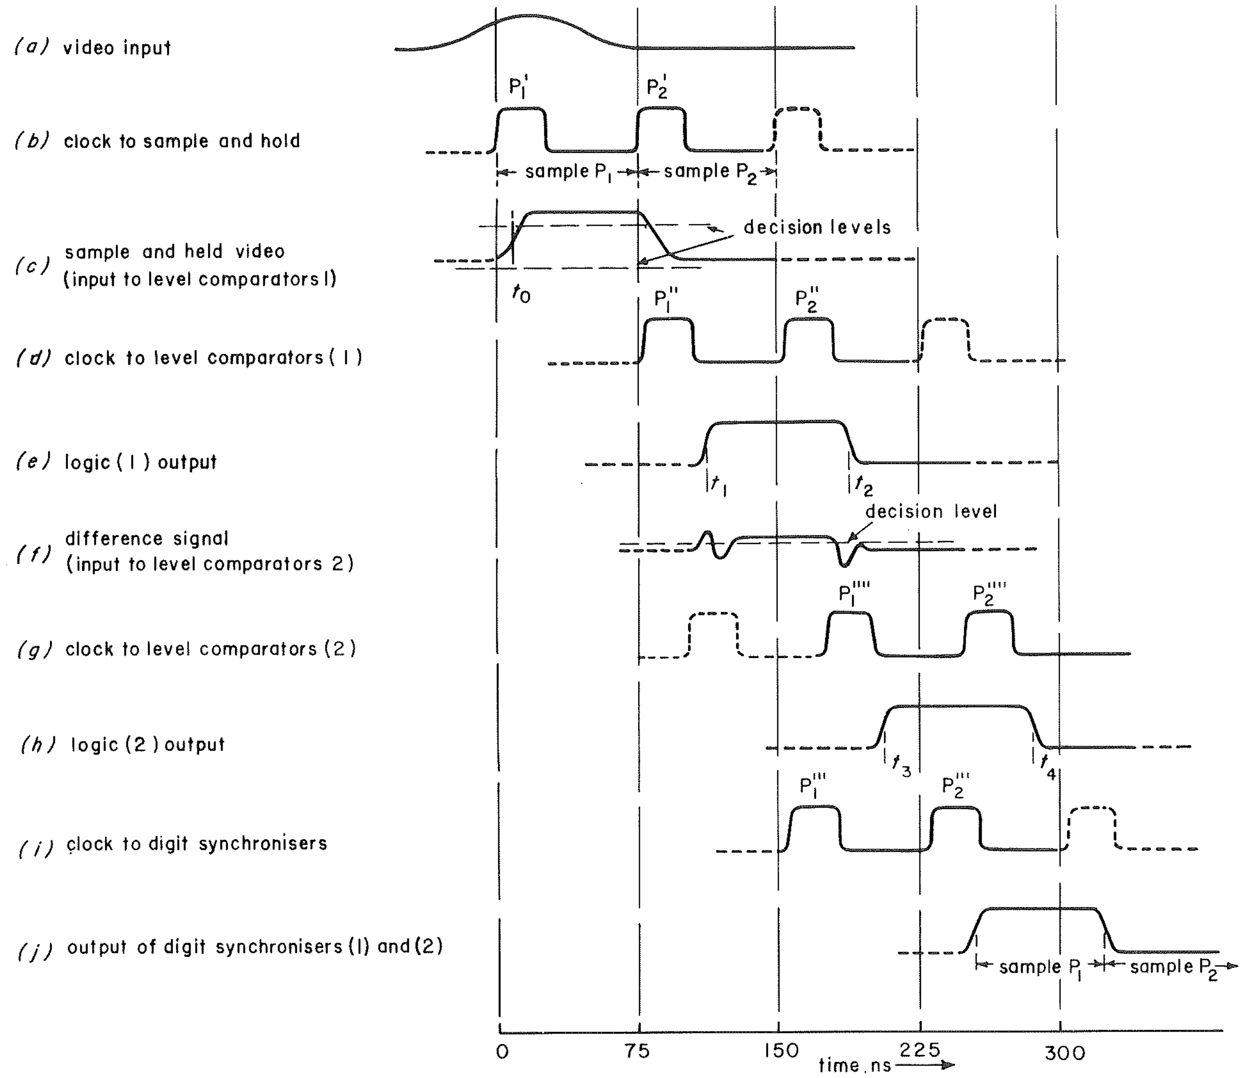

Fig. 7 - Timing sequence. Clock pulse frequency = 13·3 MHz

As previously mentioned, the two sets of digital signals corresponding to a given sample are made to occur simultaneously by passing the outputs from logic (1) through two D-type flip-flops in cascade, whereas the outputs from logic (2) are only passed through a single D-type flip-flop. The timing of the clock pulses to these flip-flops (see waveform (*i*)) is determined by the following conditions:—

(a) The leading edge of pulse P_1^{111} must occur between times t_1 and t_2 on waveform (*e*).

(b) The leading edge of pulse P_2^{111} must occur between times t_3 and t_4 on waveform (*h*).

With the clock pulses to the digit synchronisers timed as shown, all digital outputs occur simultaneously during the intervals indicated by waveform (*j*).

2.8. Reference supplies

In a video a.d.c. it is important that the quantum levels should be uniformly spaced but the absolute value of the voltage difference between quantum levels is not critical. Under these conditions, it is not necessary to take special precautions about the absolute stability of the three reference supplies required in the a.d.c. All that is required is to ensure that any change in one supply is accompanied by a directly proportional change in the other two supplies. In practice, this relationship is obtained by deriving the two reference voltages to the level comparators from the reference supply to the quantised signal generator as shown in Fig. 8.

In the circuit shown in Fig. 8, voltage drifts are minimised by using metal film resistors and integrated operational amplifiers.

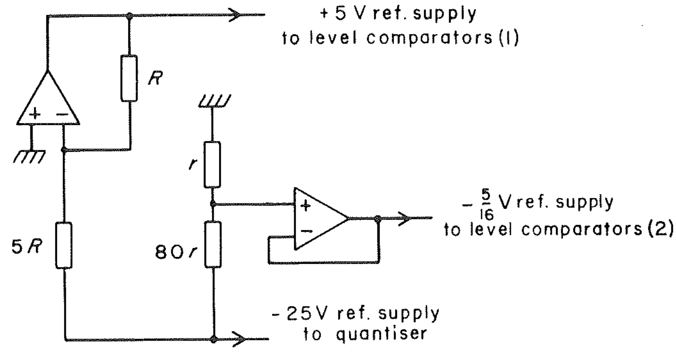

Fig. 8 - Interconnection of reference supplies

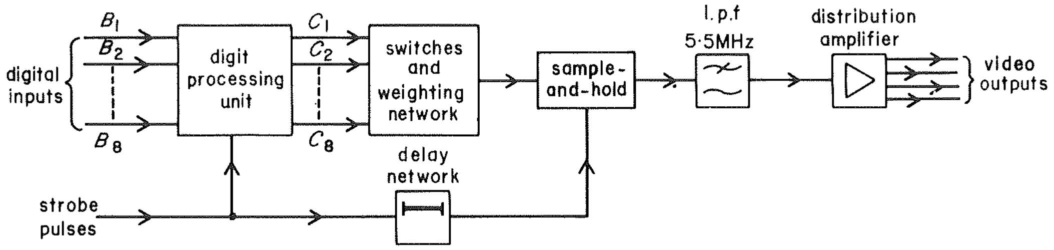

Fig. 9 - Block diagram of d.a.c.

The voltage supplies to each set of level comparators are divided into sixteen parts by a resistor chain.

2.9. Use of integrated circuits

The relatively long propagation delays allowable in each stage of the coder have made it possible for level comparator, logic and digit synchroniser stages to be constructed almost entirely from readily available integrated circuits. All the digital operations are performed with TTL integrated circuits for which the typical delay in a single NAND gate is 8 nanoseconds.

The only parts of the converter which use discrete components are those which handle analogue signals, i.e. the buffer amplifier, subtractor and sample-and-hold circuits. No linear integrated amplifiers are yet available which fulfil all the requirements of these circuits.

3. Digital-to-analogue converter

The problems involved in constructing a d.a.c. suitable for video signals have already been described elsewhere,[1,7] and therefore only an outline of the present system is given below.

The new 8-bit d.a.c. differs from the previous 6-bit d.a.c. in that weighting of the digital inputs is carried out in a resistive ladder network rather than by summing weighted currents in a single resistor.

A block diagram of the d.a.c. is given in Fig. 9. As shown in this diagram, the unit requires the digital inputs to be in parallel form and it also requires a strobe input signal consisting of a regular train of pulses occurring once per sample. In the present equipment, these strobe pulses are obtained from the a.d.c.

The principle of operation of the d.a.c. is as follows:

Firstly, skew errors (minor timing differences between nominally coincident digits) are removed by applying the digital input signals to D-type flip-flops triggered by strobe pulses. The re-timed signals from these flip-flops are then applied to switches which control currents flowing into or out of a binary resistive ladder network. This method of weighting digital signals is described in detail in reference[6], but for convenience the main circuit details are shown in Fig. 10.

Despite the re-timing of the digital signals, the quantised signal obtained from the weighting network still contains switching spikes between samples. These spikes are removed by re-sampling the quantised signal with delayed strobe pulses timed to occur approximately mid-way between switching operations.

Finally, the output signal from the re-sampler is passed through a 5·5 MHz low-pass filter, and then through an equalising network which compensates for the small inherent attenuation at high signal frequencies caused by the sample and hold process.

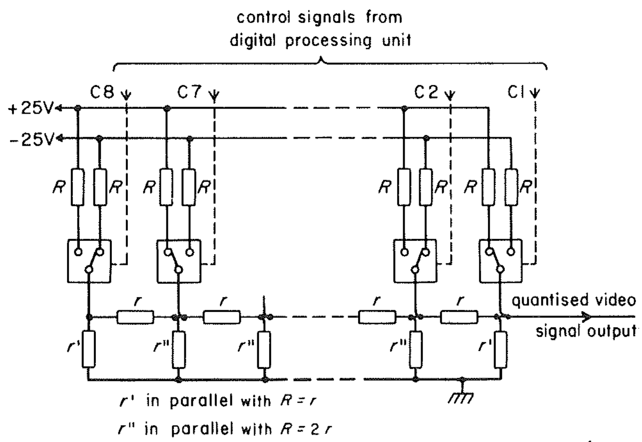

Fig. 10 - Switches and weighting network for d.a.c.

4. Mechanical construction

Both the a.d.c. and d.a.c. are constructed on plug-in units fitting into a 19 in. panel. All components are mounted on printed circuit boards. A photograph of these units together with a 5·5 MHz input filter unit is shown in Fig. 11. In the a.d.c., the sample and hold unit is mounted on the right hand vertical board, the horizontal board contains the circuits for obtaining the four most significant digits, the quantised signal generator and the subtractor, and the left hand vertical board contains the circuits for obtaining the four least significant digits.

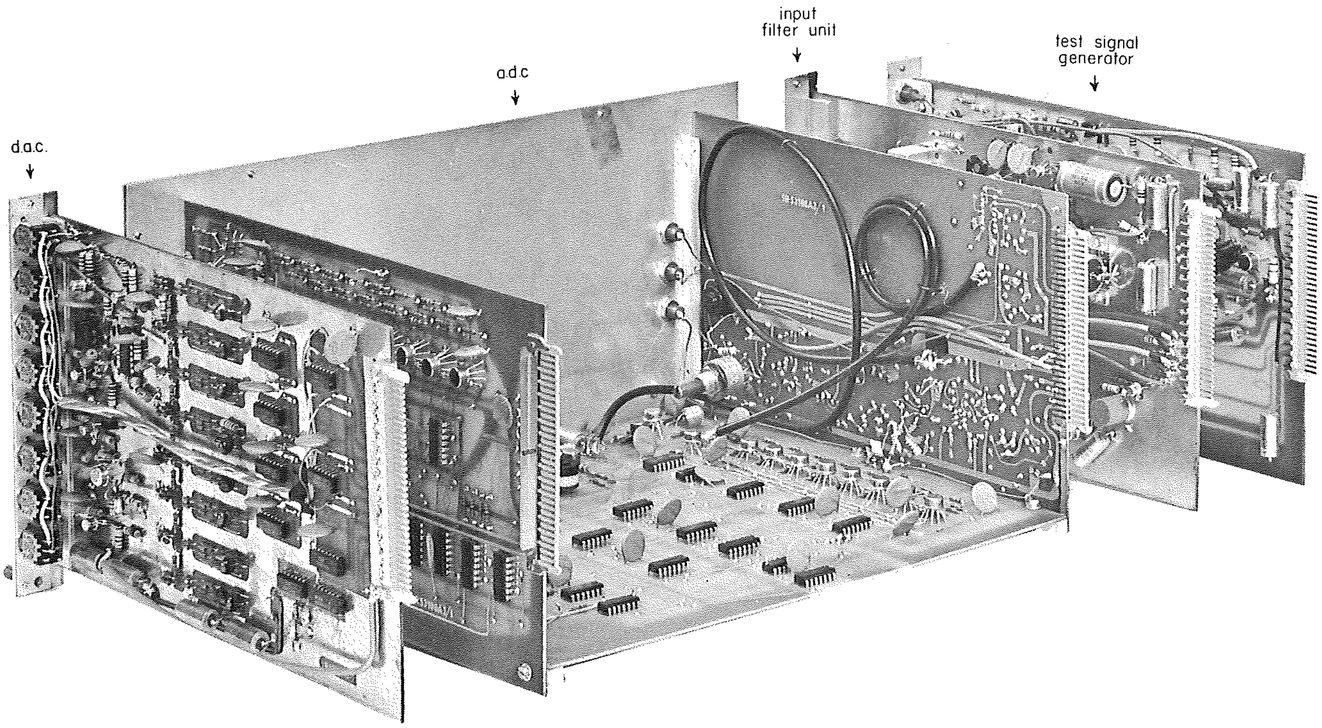

Fig. 11 - Rear view of d.a.c., a.d.c. and input filter unit

5. Performance

5.1. Accuracy of quantisation

The a.d.c. and d.a.c. were initially adjusted to make the quantum steps in the output signal as uniformly spaced as possible when a linear line-frequency ramp signal was applied to the a.d.c. The results obtained are illustrated in Fig. 12.

To obtain a quantitative figure for the accuracy of the combined coding and decoding processes, measurements were made of the peak-to-peak signal to r.m.s. quantising noise ratio obtained with both random noise and sinusoidal input signals.

The measurement techniques and theoretical predictions for a perfect system are discussed in the Appendix 2. The results obtained for the 8-bit system are given below:

For a random white noise input signal, the peak-to-peak signal to quantising noise ratio was found to be 57 dB as opposed to a theoretical figure of 59 dB.

For a 4 MHz sinusoidal input signal, the quantising noise depended on the amplitude of the signal. When the peak-to-peak magnitude of the input signal was less than about 0·2 of the conversion range, the quantising noise was 2 dB greater than the theoretical value; increasing the peak-to-peak magnitude of the signal to 0·6 of the conversion range increased the quantising noise to 4 dB greater than the theoretical figure. The main reason for this loss in conversion accuracy at high signal levels is that the difference signal from the subtractor does not settle quickly enough to its correct value after there has been a large change in signal level between successive samples.

For a 500 Hz sinusoidal input signal excursing over the full conversion range, the quantising noise was within 1 dB of the theoretical value.

These errors given above are placed in perspective by the fact that quantising noise theoretically changes by 6 dB per bit.

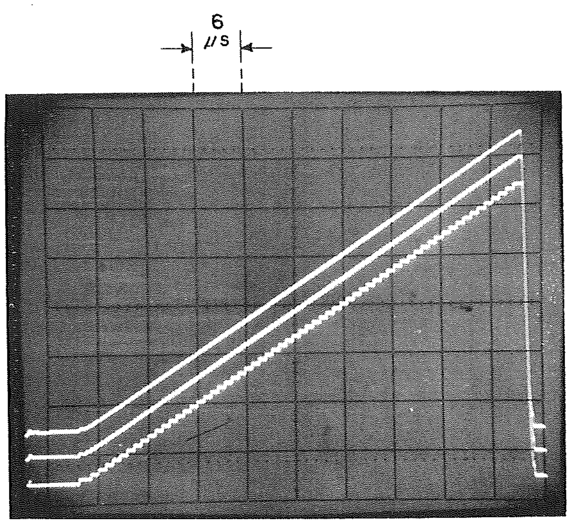

Fig. 12 - Quantised ramp signal using 6 bits (bottom trace), 7 bits (middle trace) and 8 bits (top trace)

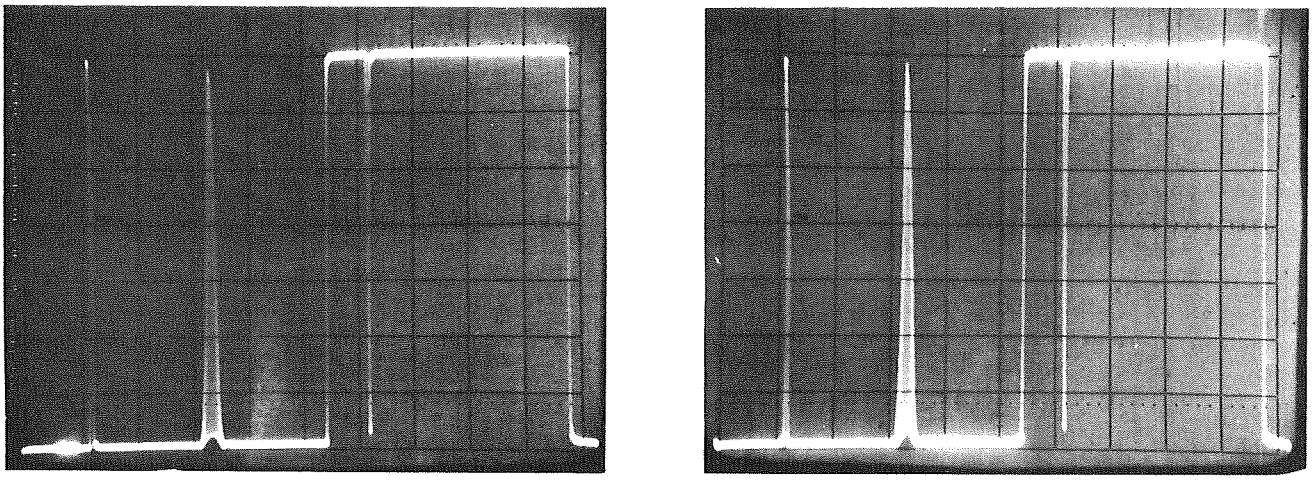

(a) (b)

Fig. 13 - Augmented pulse and bar response of system

(*a*) Input signal (*b*) Output signal

The tests with the sinusoidal input signals show that instrumental errors are negligible for slowly varying signals but improvements are desirable at high video frequencies especially for large amplitude signals. Good performance at high frequencies is particularly necessary when the video signal includes a colour subcarrier.

Using a vectorscope display, it was found that the quantising noise obtained with 100% saturated colour bars caused peak-to-peak variations of about ± 0·5° in the phase and ± 0·5% in the amplitude of the colour subcarrier.

One method of examining the response of the system to video signals was to apply an augmented pulse and bar test waveform. The results obtained are shown in Figs. 13 and the pulse response is shown on an expanded scale in Fig. 14. The slight differences between the input and output waveforms in these two figures are almost entirely caused by the two sharp cut-off 5·5 MHz low-pass filters included in the system.

5.2. Subjective effects of quantisation

These effects are to be the subject of a later report and are only described briefly opposite.

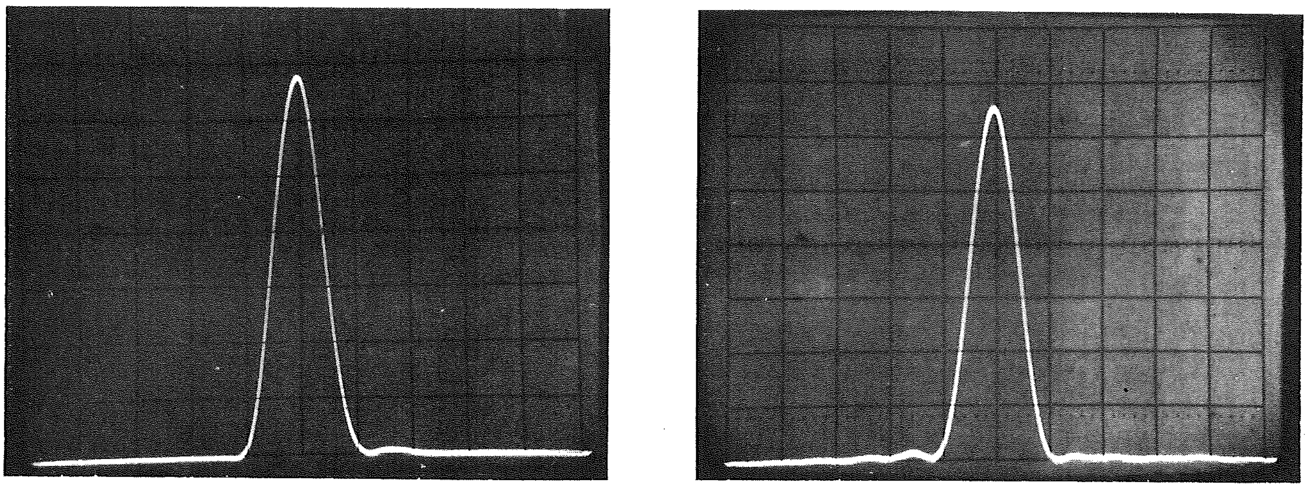

(a) (b)

Fig. 14 - 2T response of system

(*a*) Input signal (*b*) Output signal

The number of bits required if quantisation is to have a negligible effect on the picture quality depends on the type of picture being displayed. For normal programme material, it is extremely difficult to detect any impairment with seven or more bits. (In this section, it is assumed that the complete video signal including synchronising pulses has been coded.)

In the limited subjective tests so far carried out, the only types of picture on which any impairment has been noticed using 8 bits have been those obtained from test signal generators. With 8 bits, the magnitude of the noise produced by ordinary picture sources is normally sufficient to break up the unwanted effects caused by quantisation.

When white noise was added to critical types of test signal the effects of 8-bit quantisation were completely masked if the signal-to-noise ratio at the input to the a.d.c. was worse than 54 dB unweighted.

5.3. Use of dither signals to improve performance

If necessary, the subjective impairment caused by quantisation can be reduced by adding dither[8,9,10] signals to the video input of the a.d.c. These dither signals provide just sufficient high frequency information to prevent the signal from remaining between adjacent quantum levels over large areas of a picture.

A fuller discussion of the effects of dither signals will be given in a later report after more detailed tests have been performed.

6. Conclusions

The equipment described in this report shows that it is possible to construct an 8-bit video a.d.c. and and d.a.c. in which a high proportion of the circuitry is provided by readily available integrated circuits.

For a random input signal the measured peak-to-peak signal to r.m.s. quantising noise is 57 dB for 8 bits as compared with a theoretical value of 59 dB. Work is still in progress to obtain further improvements in performance.

7. References

1. Experimental equipment for converting television signals into digital signals and vice versa. BBC Research Department Report No. EL-14, Serial No. 1968/6.

2. EDSON, J.O. *and* HENNING, H.H. 1965. Broadband codecs for an experimental 224 Mb/s PCM terminal. *Bell Syst. tech. J.*, 1965, **XLIV**, 9, pp. 1887-1940.

3. TEESDALE, R.R. *and* WESTON, J.D. 1967. The encoding of broadband signals (In 2 parts). *Syst. and Commun. (G.B.)*, 1967, **3**, 2, pp. 18-22, pp. 28-30.

4. HANKE, G. 1969. PCM system in integrated circuit technique for television transmissions. (In German). *Nachrtech. Z.*, 1969, **22**, 11, pp. 621-627.

5. JOHANNESEN, F.G. 1970. Successive approximation and feedback p.c.m. encoder for television. *Proc. Inst. elect. Engrs.* 1970, **117**, 4, pp. 671-680.

6. GRAY, J.R. *and* KITSOPOULUS, S.C. 1963. A precision sample and hold circuit with sub-nanosecond switching. *IEEE Trans. Circuit Theory*, 1964, **CT-11**, 3, pp. 389-396.

7. KOVANIC, E.F. 1963. A high-accuracy 9-bit digital-to-analog converter operating at 12 Mc. *IEEE Trans. Commun. & Electron.*, 1964, No. 71, pp. 185-191.

8. LIMB, J.O. 1969. Design of dither waveforms for quantised visual signals. *Bell Syst. tech. J.*, 1969, **48**, 7, pp. 2555-2582.

9. ROBERTS, L.G. 1962. Picture coding using psuedo-random noise. *I.R.E. Trans. Inf. Theory*, 1962, **IT-8**, 2, pp. 145-154.

10. THOMPSON, J.E. 1968. Reduction of perceptual redundancy for data compression of television signals. Ph.D. thesis Imperial College, London, May 1968.

APPENDIX I

Switching levels of comparators and details of logic stages

In practice each level comparator stage in the a.d.c. contains eight pairs of differential amplifiers. A voltage difference of 2 mV between the inputs of one of these amplifiers causes the output to enter a limiting state of about 0V or +4V. The two outputs of each pair of amplifiers are connected together and fed to a single D-type flip-flop as shown in Fig. 15(*a*). This figure also shows the input signals applied to the amplifiers. The circuit design of these amplifiers is such that the output level of the pair of amplifiers is equal to the more positive level of the two separate outputs. The resulting input/output voltage characteristic is shown in Fig. 15(*b*).

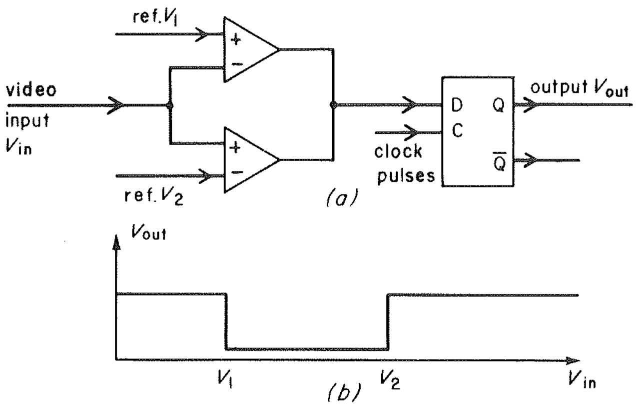

Fig. 15 - Dual level comparators
(*a*) Circuit diagram
(*b*) Voltage transfer characteristic

The relative input levels at which the eight pairs of comparators in one stage of the a.d.c. switch is indicated by waveforms A B C D E F G and H in Fig. 13.

In order to obtain an inverted quantised signal at the subtractor, the outputs obtained from logic stage 1 indicate the complement of the binary signal representing the magnitude of the video signal. The state of these output signals for varying video input levels is indicated by waveforms N_1, N_2, N_3, and N_4 in Fig. 16. These signals are inverted in digit synchroniser 1 to obtain the correct signals at the output of the a.d.c.

The logic expressions required to obtain waveforms N_1, N_2, N_3 and N_4 from the comparator outputs can be deduced from Fig. 16 and are as given below for a positive logic system:–

$N_1 = A$
$N_2 = AE + \bar{A}\,\bar{E}$
$N_3 = A\,G + \bar{A}\,\bar{G} + C\,\bar{E} + \bar{C}\,E$
$N_4 = A\,H + \bar{A}\,\bar{H} + B\,\bar{C} + \bar{B}\,C + D\,\bar{E} + \bar{D}\,E + F\,\bar{G} + \bar{F}\,G$

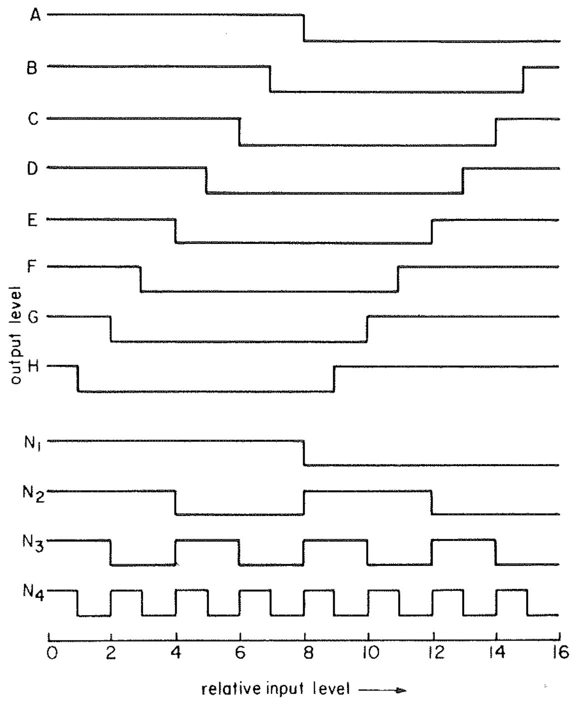

Fig. 16 - Switching levels of comparators and logic

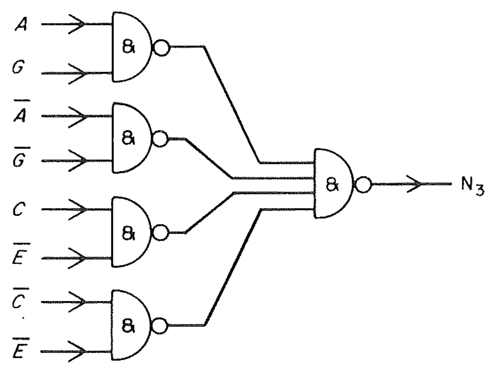

Fig. 17 - Logic circuitry to obtain 3rd most significant bit

As an example of the logic circuits required, the method of obtaining waveform N_3 is shown in Fig. 17.

Logic stage 2 is identical to logic stage 1. No inversion of the binary signals is needed in digit synchroniser 2, as the video input to the second stage is inverted with respect to the video input of the first stage.

APPENDIX II

Measurement of quantising noise

The errors introduced by a quantisation process are often referred to as quantising noise. For a perfect linear quantiser with a random input signal, the ratio of the maximum voltage swing handled by the quantiser to the r.m.s. quantising noise is equal to $(6n + 11)$ dB where n is the number of bits. The accuracy of quantisation of a system can be measured by comparing the actual value of the quantising noise with this theoretical value.

It should be noted that if the entire video signal including synchronising pulses and 100% saturated colours is coded, then the video signal-to-quantising noise ratio is equal to $(6n + 6)$ dB. This is because the black-to-white excursion of a video signal is used as a reference in video signal-to-noise measurements and this signal is 5 dB less than the maximum possible swing in a composite colour video signal.

To measure quantising noise it must first be separated from the wanted component of the quantised signal. For the purpose of the present work this separation was carried out by filtering techniques which are described below for different types of input signal.

(a) 500 Hz sinusoidal input signal

The 500 Hz component was removed from the quantised signal by means of a 10 kHz high-pass filter. This filter also removes the quantising noise in the range 0 – 10 kHz but the resulting error is small for signals crossing a large number of quantum levels and it has been ignored. A comparatively high value of cut-off frequency was chosen to prevent harmonics of 500 Hz present in the input signals being measured as quantising noise.

(b) 4 MHz sinusoidal input signal

Harmonics of 4 MHz in the input signal were removed by the 5·5 MHz low-pass filter normally used at the input of the a.d.c. when video signals are being coded. The 4 MHz component in the quantised signal was removed by means of a 4 MHz band-stop filter having a bandwidth of 200 KHz.

(c) Random white noise input signal

The spectrum of the input noise signal was flat over the video band apart from a slot filtered out by a 4 MHz band-stop filter having a bandwidth of 200 kHz. Since quantising noise has a flat amplitude spectrum, the total quantising noise power can be calculated by measuring the noise power introduced into this slot. For this measurement, the quantised signal was passed through a 4 MHz band-pass filter with a bandwidth of 50 kHz.

Printed by BBC Research Department, Kingswood Warren, Tadworth, Surrey